

1. Which category of machine learning algorithms is primarily concerned with learning through trial and error, aiming to maximize cumulative reward?

- a) Supervised learning
- b) Unsupervised learning
- c) Reinforcement learning
- d) Semi-supervised learning

Answer: c) Reinforcement learning

Reinforcement learning (RL) involves an agent interacting with an environment to learn a policy that maximizes cumulative rewards.

2. Which bandit algorithm aims to balance exploration and exploitation by using Upper Confidence Bound (UCB)?

- a) ϵ -greedy
- b) Thompson Sampling
- c) PAC (Probably Approximately Correct)
- d) UCB

Answer: d) UCB

UCB algorithms estimate the upper confidence bounds for the expected rewards of each action, guiding the agent to explore less explored actions while exploiting actions with higher

expected rewards.

3. Which bandit algorithm guarantees to find the optimal arm with high probability within a finite number of steps?

- a) UCB
- b) PAC
- c) Median Elimination
- d) Thompson Sampling

Answer: c) Median Elimination

Median Elimination algorithm is designed to eliminate suboptimal arms quickly and converge to the optimal arm with high probability within a finite number of steps.

4. Which RL method directly learns the policy by optimizing the expected cumulative reward through gradient ascent?

- a) Policy Gradient
- b) Q-learning
- c) Dynamic Programming
- d) Temporal-Difference Learning

Answer: a) Policy Gradient

Policy Gradient methods aim to directly optimize the policy parameters by maximizing the expected cumulative reward through gradient ascent.

5. Which RL concept involves modeling the environment as a Markov Decision Process (MDP) to formulate learning problems?

- a) Bellman Optimality
- b) Temporal-Difference Learning
- c) Function Approximation
- d) Eligibility Traces

Answer: a) Bellman Optimality

Bellman Optimality principle is a foundational concept in RL that helps in formulating optimal policies within the context of MDPs.

6. Which dynamic programming method involves iteratively improving the value function by applying the Bellman Optimality equation?

- a) Value Iteration

- b) Policy Iteration
- c) Q-learning
- d) Temporal Difference Learning

Answer: a) Value Iteration

Value Iteration iteratively improves the value function by applying the Bellman Optimality equation until convergence.

7. In which dynamic programming method does the algorithm alternate between policy evaluation and policy improvement steps until convergence?

- a) Value Iteration
- b) Policy Iteration
- c) Q-learning
- d) Temporal Difference Learning

Answer: b) Policy Iteration

Policy Iteration alternates between evaluating the current policy and improving it until an optimal policy is found.

8. Which RL algorithm combines elements of dynamic programming and function approximation, particularly in large state spaces?

- a) Value Iteration
- b) Policy Iteration
- c) Q-learning
- d) Function Approximation

Answer: d) Function Approximation

Function Approximation techniques are often used in RL to deal with large state spaces by approximating value functions or policies.

9. Which RL method updates value estimates by bootstrapping from the current estimate towards a target that includes the reward plus the estimated value of the next state?

- a) Value Iteration
- b) Policy Iteration
- c) Q-learning
- d) Temporal Difference Learning

Answer: d) Temporal Difference Learning

Temporal Difference Learning updates value estimates based on the difference between current and future estimates, integrating both immediate rewards and future predictions.

10. Which technique in RL is used to credit or blame actions taken by the agent based on their influence on the received reward?

- a) Temporal Difference Learning
- b) Eligibility Traces
- c) Function Approximation
- d) Least Squares Methods

Answer: b) Eligibility Traces

Eligibility Traces attribute credit or blame to actions based on their influence on the received reward, aiding in updating value estimates efficiently.

11. Which method in RL involves approximating value functions or policies using linear regression to handle large state spaces?

- a) Temporal Difference Learning
- b) Function Approximation
- c) Q-learning

d) Least Squares Methods

Answer: d) Least Squares Methods

Least Squares Methods use linear regression to approximate value functions or policies, making them suitable for handling large state spaces.

12. In RL, which technique is used to update the action-value function towards the target value computed using a bootstrapped estimate?

- a) Value Iteration
- b) Policy Iteration
- c) Q-learning
- d) Temporal Difference Learning

Answer: c) Q-learning

Q-learning updates the action-value function towards the target value computed using a bootstrapped estimate, facilitating learning optimal policies.

13. Which RL method aims to directly optimize the policy without explicitly computing the value function?

- a) Policy Gradient
- b) Q-learning
- c) Value Iteration
- d) Policy Iteration

Answer: a) Policy Gradient

Policy Gradient methods aim to directly optimize the policy without requiring explicit computation of the value function.

14. Which bandit algorithm chooses actions probabilistically based on posterior distributions of each arm's reward?

- a) ϵ -greedy
- b) Thompson Sampling
- c) PAC
- d) UCB

Answer: b) Thompson Sampling

Thompson Sampling selects actions probabilistically based on posterior distributions, making decisions that balance exploration and exploitation.

15. Which bandit algorithm guarantees a high probability of selecting the optimal arm within a specified number of rounds?

- a) ϵ -greedy
- b) Thompson Sampling
- c) PAC
- d) UCB

Answer: c) PAC

PAC (Probably Approximately Correct) guarantees a high probability of selecting the optimal arm within a specified number of rounds.

16. Which RL technique aims to approximate the value function or policy using a neural network?

- a) Q-learning
- b) Function Approximation
- c) Policy Gradient
- d) Dynamic Programming

Answer: b) Function Approximation

Function Approximation techniques use neural networks to approximate value functions or policies, enabling RL in complex environments.

17. Which RL algorithm combines elements of dynamic programming with exploration-exploitation strategies through the use of an ϵ -greedy policy?

- a) Value Iteration
- b) Q-learning
- c) Policy Gradient
- d) Policy Iteration

Answer: b) Q-learning

Q-learning combines dynamic programming principles with exploration-exploitation strategies through ϵ -greedy policy, facilitating learning optimal policies.

18. Which bandit algorithm balances exploration and exploitation by choosing a random action with probability ϵ and the best known action otherwise?

- a) Thompson Sampling
- b) UCB
- c) Median Elimination

d) ϵ -greedy

Answer: d) ϵ -greedy

ϵ -greedy algorithm balances exploration and exploitation by choosing random actions with probability ϵ and the best known action otherwise.

19. Which RL method estimates the value of each state-action pair and updates these estimates based on the difference between predicted and actual rewards?

- a) Value Iteration
- b) Policy Iteration
- c) Q-learning
- d) Temporal Difference Learning

Answer: c) Q-learning

Q-learning estimates the value of each

state-action pair and updates these estimates based on the difference between predicted and actual rewards, facilitating policy improvement.

20. In which RL method does the agent learn by interacting with the environment and adjusting its policy to maximize cumulative rewards over time?

- a) Policy Gradient
- b) Q-learning
- c) Value Iteration
- d) Policy Iteration

Answer: b) Q-learning

Q-learning allows the agent to learn by interacting with the environment, adjusting its policy to maximize cumulative rewards over time through iterative updates.

21. Which RL algorithm is used to solve Markov Decision Processes (MDPs) with finite state and action spaces by iteratively improving value estimates?

- a) Value Iteration
- b) Policy Iteration
- c) Temporal Difference Learning
- d) Policy Gradient

Answer: a) Value Iteration

Value Iteration is used to solve MDPs with finite state and action spaces by iteratively improving value estimates until convergence.

22. Which bandit algorithm eliminates suboptimal arms in each round and focuses on exploring promising arms?

- a) Thompson Sampling
- b) UCB
- c) Median Elimination
- d) ϵ -greedy

Answer: c) Median Elimination

Median Elimination algorithm eliminates suboptimal arms in each round and focuses on exploring promising arms, converging to the optimal arm with high probability.

23. In RL, which technique is used to approximate the value function or policy using a linear combination of features?

- a) Temporal Difference Learning
- b) Function Approximation
- c) Policy Gradient

d) Eligibility Traces

Answer: b) Function Approximation

Function Approximation approximates the value function or policy by combining features linearly, allowing RL in large state spaces.

24. Which RL method directly updates the policy parameter by considering the gradient of the expected cumulative reward?

- a) Q-learning
- b) Policy Gradient
- c) Value Iteration
- d) Policy Iteration

Answer: b) Policy Gradient

Policy Gradient directly updates the policy parameter by considering the gradient of the expected cumulative reward, facilitating policy optimization.

25. Which bandit algorithm estimates the upper confidence bounds for each action and selects the action with the highest upper confidence bound?

- a) Thompson Sampling
- b) UCB
- c) Median Elimination
- d) ϵ -greedy

Answer: b) UCB

UCB estimates the upper confidence bounds for each action and selects the action with the highest upper confidence bound to balance exploration and exploitation.

26. Which dynamic programming method involves evaluating the current policy and improving it iteratively until convergence?

- a) Value Iteration
- b) Policy Iteration
- c) Q-learning
- d) Temporal Difference Learning

Answer: b) Policy Iteration

Policy Iteration involves evaluating the current policy and improving it iteratively until convergence, aiming to find an optimal policy.

27. Which RL method combines elements of dynamic programming with function approximation, particularly in continuous state spaces?

- a) Q-learning
- b) Function Approximation
- c) Policy Gradient
- d) Temporal Difference Learning

Answer: b) Function Approximation

Function Approximation combines dynamic programming with function approximation, making it suitable for continuous state spaces in RL.

28. Which RL concept involves updating value estimates based on the difference between successive estimates and the reward received?

- a) Bellman Optimality
- b) Dynamic Programming
- c) Temporal Difference Learning
- d) Policy Iteration

Answer: c) Temporal Difference Learning

Temporal Difference Learning updates value estimates based on the difference between successive estimates and the reward received, facilitating learning without requiring a model of the environment.

29. In RL, which method involves iteratively improving the value estimates by applying the Bellman Optimality equation and maximizing cumulative rewards?

- a) Policy Gradient
- b) Value Iteration
- c) Q-learning
- d) Eligibility Traces

Answer: b) Value Iteration

Value Iteration iteratively improves value estimates by applying the Bellman Optimality equation and maximizing cumulative rewards, facilitating optimal policy determination.

30. Which bandit algorithm aims to eliminate arms with estimated mean below a certain threshold in each round?

- a) UCB
- b) Thompson Sampling

- c) Median Elimination
- d) ϵ -greedy

Answer: c) Median Elimination

Median Elimination aims to eliminate arms with estimated mean below a certain threshold in each round, focusing exploration on promising arms.

Related posts:

1. Introduction to Information Security
2. Introduction to Information Security MCQ
3. Introduction to Information Security MCQ
4. Symmetric Key Cryptography MCQ
5. Asymmetric Key Cryptography MCQ
6. Authentication & Integrity MCQ
7. E-mail, IP and Web Security MCQ